An Introduction to IPv6

By Brian Tamburello

The Internet, as well as most mid-to-large networks, is built on a system known as a TCP/IP stack (Transmission Control Protocol/Internet Protocol). TCP/IP was initially developed by the U.S. Department of Defense (DoD) in the late 1960s.[1] Internet Protocol works at the networking layer of this system, and as the Internet has developed, it has gone through several adaptive evolutions to maintain relevant functionality. The most widely used version of this protocol (version 4 or IPv4) is nearly 20 years old, and will, according to most sources, gradually be replaced by a newer protocol: version 6 or IPv6.

IPv6, sometimes called the Next Generation Internet Protocol or IPng, was developed by the Internet Engineering Task Force (IETF) in the early to mid-1990s.[2] IPv6 improvements over IPv4 include: (1) better scalability (due to more addressable space); (2) new security specifications; (3) better support for higher bandwidth/real-time traffic; (4) plug-and-play networking; (5) easy device renumbering; and (6) an automatic address clustering system. Additionally, it builds upon and improves the good optimization and logical structure of IPv4 while breaking away from some of its shortcomings.[3]

Replacing IPv4

Despite its age, IPv4 has shown considerable resiliency and has been extended in many ways to meet the needs of the rapidly growing Internet. There is; however, no easy solution for its most pressing problem: IPv4 uses a 32-bit standard. Under this system, every networked device requires its own unique IP address, usually denoted in a "dotted quad" decimal format, for example, 192.168.1.1.[1] In other words, there are, theoretically, 2^{32} individual addressable numbers available for use. In reality, the usable address space is somewhat less than that because there are many number ranges reserved for special purposes and network hierarchies. When the standard was initially developed, 32 bits seemed more than sufficient to address every computer worldwide, but more recent estimates have shown that at the current rate of IP address consumption, IPv4 will be out of addressable space within the next 10 years.[3]

Among experts, there is some controversy regarding what the IPv4 replacement standard will actually be. There are some who think that IPv4's address shortage problem can be overcome through judicious use of "Network Address Translation" (NAT). Others think that a complete redesign of Internet Protocols are needed to achieve the desired result and some unrelated standard will replace IPv4. The next logical replacement would be a 64-bit protocol, and one was created, IPv5, but it existed solely in an experimental form. The IPv5 header served to identify an experimental streaming packet protocol called ST. ST never gained widespread industry acceptance, but since the header was already used, the next obvious choice was a 128-bit protocol, IPv6.[2]

IPv6 Addressing

The IPv6 specification calls for a 128-bit address, which theoretically provides 2^{128} or 40 undecillion (40*10^66) unique addresses, yielding 1000 addresses for every square meter of surface area on Earth.[4] In practice that much space won't be available due to various network inefficiencies and reserved value ranges, but there are still clearly a very large number of addressable values available. Unlike the IPv4 decimal-based system, IPv6 uses hexadecimal values broken up into 8 groups of 4 values per group in the following manner:

0123:4567:89ab:cdef:0123:4567:89ab:cdef

For the sake of convenience and efficiency any leading zeros may be dropped, so:

fa36:004d:0000:0000:0000:0000:34bb:0001

Could be truncated to:

fa36:4d:0:0:0:0:34bb:1

Another format that allows further efficiency (referred to as "compressed" format) allows long sequences of zeros to be denoted simply as a double colon ("::"). The above address could then be written as follows: fa36:4d::34bb:1

These features will be especially useful in the earlier phases of address issuing because addresses containing "low" values (containing many zeros) will be quite common.

A third, final format is specified for a mixed IPv4/IPv6 environment (referred to as "IPv4 compatible") and allows the last 32-bits of the IPv6 address to store the IPv4 address (in decimal format). The first 96-bits are used to store a hexadecimal prefix identifying the host and signaling the existence of an IPv4 compatible address.[4] For instance, if you want to store an IPv4 address, such as, 192.168.1.1, in an IPv6 format, it could be written as follows:

0:0:0:0:0:ffff:192.168.1.1

This format also supports the compressed scheme:

::ffff:192.168.1.1

Address Autoconfiguration

One of the most important features of IPv6 to the end user is its automatic address configuration. To gain access to the Internet using IPv4, one must either manually set the IP address, network mask and default gateway or connect through a Dynamic Host Control Protocol (DHCP) server, which allows for automated IP address assignment. While IPv6 does allow manual configurations, it can be handled automatically with no reliance on outside systems using its "Stateless Address Autoconfiguration Protocol." This is accomplished by setting the right-most 64-bits of the IPv6 address to be the host's Interface ID, which is based on the system's unique 48-bit Media Access Control (MAC) address on its Ethernet device. After generating this number, it then checks to see if a duplicate address exists on the Internet. If it does, a randomly generated value will be substituted for the MAC address. This should happen rarely because there are procedures in place to prevent issuing the same MAC addresses to multiple Ethernet units.

This safeguard was necessary not only because there are some legitimate uses of having identically addressed units, but also because Ethernet manufacturers have been known to accidentally issue duplicate MAC addresses. At this point, the system will then check to see if it can access an IPv6-capable router and, if so, will generate the left-most 64 bits of the address based on the ID of the host subnet and its public topology routing prefix, which specifies the location of the router.[2]

Security

The Internet currently has a number of security issues. The first issue is that all of its major privacy and authentication systems are implemented at the application level. IPv6 offers two integrated, low-level security features, which may be used individually or in tandem. The first mechanism is the "IPv6 Authentication Header," which is used to authenticate data integrity and sources, but does not, by itself, assure user confidentiality. The header supports multiple data authentication algorithms, but requires a keyed "MD5" (a popular authentication algorithm) validation routine to verify users/sources. Not only does having a standard for user authentication help assure interoperability, but it will also help prevent relatively common network attacks based on host masquerading. By providing this service at the low Internet level, applications that do not currently have their own host and data validation schemes can take advantage of this feature fairly easily.

The second integrated security feature is the "IPv6 Encapsulating Security Header," which allows user/host confidentiality. Like the Authentication Header, it is algorithm independent, but to encourage interoperability the "standard" algorithm currently used is known as DES CBC.5. When used together, these security protocols allow for anonymous but verifiable secure data transfers.

The "6Bone" Network

The IPv6 backbone or "6bone," is an international, experimental testing ground for IPv6 technology. This worldwide network, which is actually composed of several smaller, regional 6bone networks, is run by the IETF IPng Transition Working Group. This group's responsibilities include: (1) testing the standards and implementations of IPv6; and (2) testing transition strategies and providing opportunities for Internet Service Providers (ISP), application developers and hardware manufacturers to test the protocol prior to rolling out their own products and services.

Since most ISPs and private IPv4 communication providers have not yet integrated IPv6 routers into their systems, many IPv6 testers are forced to connect to other parts of the 6bone network through the existing IPv4 infrastructure. One of IPv6's transition features, "IPv6 Encapsulation," is particularly useful in this case. It allows the protocol to tunnel through the IPv4 routers to access the 6bone network unimpeded.[2]

Transition

The most important aspect of the transition to IPv6 involves getting it to coexist with existing IPv4 hosts, and accomplishing this well before the older standard reaches its maximum capacity. A second, related goal is to allow the transition to IPv6 hardware and software to occur in an incremental and diffusible fashion. Upgrading all the networks in the world at once or in any particular order would be difficult, if not impossible. Allowing any part of any network to be updated at any time while still maintaining a seamless system prevents this from becoming an issue, and it is critical to short-term IPv6 acceptance.

A third, less crucial objective is for the transition process to be as painless and easy as possible for end users and network operators. If the process advances as hoped, the changes would be virtually transparent to most users. IPv6 possesses transition friendly protocols making these goals achievable; it also has the invaluable advantage of requiring very little work and money to upgrade existing IPv4 systems to the new standard.[5] In fact, a great deal of existing and soon-to-be released software has been created with IPv6 compatibility specifically in mind. Microsoft, for instance, has already implemented a version of IPv6 into Windows XP and .NET Server. Although it is not activated by default, a simple command line option can enable it.[1]

Conclusions

It is obvious that IPv4, in its current form, is rapidly approaching the end of its ability to expand and adapt. Although it may appear that a global transition to IPv6 is an inevitable step that the ever-evolving Internet must take, it is very possible some other path might be taken that better suits its needs in coming years since there are other competing standards for the next generation of networking systems.

IPv6 is simply the official successor to the system currently in use, but because ease of transition was a key element in its design, it has a critical advantage over its competitors and currently has a great deal of industry support.[5] Unless some viable, long-term extension of IPv4 is created in the near future, IPv6 (or a combination of the two) will be the path of least resistance to the next generation of networking services. It is a realistic, evolutionary step that builds on all IPv4's successes and strengths. IPv6 will eliminate all of IPv4's weaknesses, and it will be able to offer many revolutionary networking advancements.

References

1. Shinder, Deb. "IPv6: What is it and why is it needed?" *TechRepublic, May 2002.* (http://www.zdnet.com.au/insight/0,39023731,20265559,00.htm)

2. Fink, Robert L. "IPv6 - What and Where it is." *The Internet Protocol Journal, March 1999.* (http://www.cisco.com/warp/public/759/ipj_2-1/ipj_2-1_ipv6.html)

3. WIDE Project, *IPv6 Working Group. Frequently Asked Questions.* (http://www.itojun.org/v6/v6faq.html)

4. Scholz Gregory R. "Internet Protocol Version 6." *Consortium for Computing in Small Colleges, 2001.* (http://portal.acm.org/citation.cfm?id=374779&jmp=abstract&dl=GUIDE&dl=ACM)

5. Hinden Robert M. "IP Next Generation Overview." *Communications of the ACM, June 1996.* (http://portal.acm.org/citation.cfm?id=228503.228517&dl=portal&dl=ACM)

Brian Tamburello is a computer scientist in the Technical Specifications and Acquisitions Branch, SPAWARSYSCEN Charleston, DON IT Umbrella Program Norfolk Office. He has a Bachelor of Science degree in computer science from Virginia Commonwealth University.